



Query-Reusable Proof Systems

Yuval Ishai^{1,2} , Eyal Kushilevitz¹ , Varun Narayanan³ ,
Rafail Ostrovsky³ , and Akash Shah³

¹ Technion - Israel Institute of Technology, Haifa, Israel

² AWS, Seattle, USA

³ University of California Los Angeles, Los Angeles, USA
varunnkv@gmail.com

Abstract. Probabilistic proof systems such as PCPs and their zero-knowledge variants (ZK-PCPs) are central building blocks in cryptographic applications. In this work, we study *query-reusable* proof systems where the verifier can sample its queries once and use them to verify any polynomial number of proofs. In this reusable setting, soundness should still hold even if the prover can learn the verifier’s decision (accept or reject) on many badly formed proofs. Our study is motivated by attractive features of designated-verifier NIZK systems that combine a query-reusable (honest-verifier) ZK-PCP with symmetric encryption.

The reusability of ZK-PCP was studied by Chase et al. (Crypto 2019), who obtained a limited negative result for ZK-PCP with a special simulator. This left the question open for unrestricted ZK-PCP. We essentially settle this question by showing a negative result for statistical ZK-PCP (alternatively, PCP with sublinear query complexity) under standard complexity theoretic assumptions. We complement this with a positive result, showing that if *either* soundness or ZK are computational, query-reusable ZK-PCPs that *do not* meet the special simulation requirement of Chase et al. follow from standard cryptographic assumptions.

Finally, we study the relaxed notion of *bounded* query reusability, where the prover is allowed to interact with the verifier over a bounded number of epochs by issuing a batch of polynomially many proofs in each epoch and learning the verifier’s decisions. We obtain a nearly tight characterization of the number of queries required for r -epoch reusability.

1 Introduction

Probabilistic proof systems such as Probabilistically Checkable Proofs [2] (PCPs) and their Zero-Knowledge variants [21] (ZK-PCPs) are central building blocks in many cryptographic applications [1, 3, 4, 6, 12, 15, 19, 23]. In this work, we revisit the notion of *query-reusable* proof systems considered in [7].

We begin with an example that motivates our work. Consider the goal of obtaining information-theoretic non-interactive zero-knowledge (NIZK) proofs, with a designated verifier, using a correlated randomness setup. Kilian, Micali and Ostrovsky describe such a solution in [20] where the prover gets a set of keys, and the verifier gets a randomly chosen subset of the keys. The prover

publishes a proof after masking it with the keys. Using the keys to unmask a random subset of the proof, the verifier is able to keep a “watchlist”, and base their decision on the consistency in the watchlist.

Can the correlation setup in such an NIZK be reused? In other words, can the verifier reuse the same watchlist across all proofs? If we settle for computational zero-knowledge, the correlation can provide seeds of a pseudo-random function (PRF) instead of one-time pads. Now, the prover can generate a fresh key for each new instance by evaluating the PRF on a new point. Clearly, the resulting scheme inherits completeness from the original scheme. Soundness also holds if a malicious prover learns no information about whether the verifier accepts or rejects the proofs. However, as observed in [20], if the prover learns the verifier’s decision, this approach to reusability encounters a subtle but common “selective failure attack”, where the prover can learn partial information about the verifier’s secret randomness by generating suitable malformed proofs. The setup in [20] uses several independent copies of a simple correlation in which verifier receives one among three possible subsets of prover’s keys. This allows a malicious prover to learn the verifier’s entire randomness from the verifier’s decision on polynomially many non-adaptively chosen malformed proofs, and break soundness. The question whether a similar attack applies to *every* proof system of this kind was studied in [7], in the broader context of reusable secure two-party computation.

Towards a general study, it is convenient to abstract the type of zero-knowledge proof systems from [20] using the notion of an (honest-verifier) *zero-knowledge PCP*, or ZK-PCP for short. In a ZK-PCP, a computationally bounded verifier samples a set of coordinates, referred to as the *query set*, according to some predetermined distribution, and decides to accept or reject after checking the proof published by the prover on the set of coordinates. In addition to ensuring completeness and soundness, a ZK-PCP ensures that the distribution of the proof restricted to the query set can be simulated using only the input. The quality of the simulation determines the type of zero-knowledge obtained.

A ZK-PCP for an NP language is query-reusable if the verifier can sample the query set once and use it to verify polynomially many proofs issued by the prover while maintaining soundness. The soundness condition requires that a malicious prover, who learns the verifier’s decision on adaptively chosen inputs and potentially malformed proofs, cannot convince the verifier to accept an input outside the language except with negligible probability.

NIZK systems that combine a reusable honest-verifier ZK-PCP with symmetric encryption have several highly attractive features. From a theory perspective, they can be based on a standard and “minimal” assumption: the existence of one-way functions. In contrast, known NIZK protocols from standard cryptographic assumptions (even in the designated-verifier setting we consider here) are all based on “structured” assumptions related to factoring [10], discrete log [8, 9, 18, 22], lattices [22, 24, 25], or learning parity with noise [22]. From a practical point of view, while the Fiat-Shamir heuristic gives rise to NIZK protocols that rely on symmetric cryptography alone, this approach has two sources of con-

crete overhead. First, cryptographic hash functions are typically more expensive than symmetric encryption. Second, and more interestingly, in order to apply the Fiat-Shamir heuristic, one inherently needs to start with an interactive protocol that has a very low soundness error, typically 2^{-128} . Achieving this level of soundness for interactive protocols based on zero-knowledge PCP (e.g., using the MPC-in-the-head paradigm [1, 6, 15]) requires repeating a baseline protocol many times for soundness amplification. In contrast, with the “watchlist” approach we employ here, one can potentially get much better efficiency by settling for a much higher soundness error that often suffices in practice. In fact, even soundness error of $1/2$ may be enough in situations where a malicious prover may suffer reputation loss or be subject to legal consequences when caught cheating. Thus, obtaining NIZK protocols based on *reusable* watchlists is a practically motivated problem.

In [7], the authors partially address this problem by showing that ZK-PCP with a strong notion of zero knowledge cannot be reused, but leave the general question open. Specifically, their result doesn’t apply to NIZK systems that use honest-verifier ZK-PCP, which served as the initial motivation for our work. Another line of work—[17, 21] and, recently, [14]—studies ZK-PCP with malicious verifier with the objective of building sublinear ZK-PCP where a malicious verifier can make many more adaptive queries than a semi-honest verifier. This line of work is less relevant to our work since we only consider semi-honest or “semi-malicious” verifiers, or alternatively non-succinct ZK-PCP with OT query pattern. Note, this makes our negative results stronger.

1.1 Our Contributions

We initiate a broad study of reusable probabilistic proof systems. The impossibility of reusability in [7] applied even for ZK-PCP with computational zero-knowledge even against computationally bounded provers, referred in the literature as ZK-PCP arguments. However, zero-knowledge property required that the distribution of the proof generated by the prover restricted to any given set of coordinates (of bounded size) can be simulated. This simulator is much stronger than one demanded by honest verifier ZK-PCP, which simulates the distribution of the proof restricted to a random query set. We address reusability of honest-verifier ZK-PCP with the standard notion of zero-knowledge. In the sequel we will only consider honest-verifier setting, hence we will suppress this prefix. We obtain the following result:

Theorem 1 (Negative result for reusable ZK-PCP, informal). *Under a standard complexity assumption, statistical ZK-PCP is not query-reusable.*

We then show that the limitation on simulation in the negative result for ZK-PCP arguments imposed in [7] is in fact necessary, in the sense that without this restriction there is a reusable PCP argument under the standard learning with errors (LWE) assumption.

Theorem 2 (Positive result for reusable ZK-PCP arguments, informal). *Assuming LWE , there exist honest-verifier perfect ZK-PCP arguments (alternatively, computational ZK-PCP proofs) that are query-reusable.*

Next, we turn to other kinds of proof systems: classical (sublinear) PCPs, Witness-Indistinguishable PCPs, and linear PCPs, and obtain similar negative results to Theorem 1. See Sect. 1.2 for an overview of them.

The attacks that establish the above negative results require many “epochs” of interaction with the query-reusing verifier. In each epoch, based on what it learned in the previous epochs, the adversary chooses a batch of polynomially many inputs and associated purported proofs, and receives the corresponding decision of the query-reusing verifier. This leads to the following question:

Can we build PCP-style proof systems that are reusable for a bounded number of epochs even if the number of proofs verified in each epoch is unbounded?

All standard (“classical”) PCPs and ZK-PCPs from the literature that we are aware of can be broken following just one epoch. For instance, the ZK-PCP implicit in [20] is built using a simple perfect ZK-PCP with soundness $1/2$ which consists of three strings out of which the verifier queries two at random. The soundness of this scheme is boosted by issuing several independently sampled ZK-PCPs in parallel. An adversarial prover can learn the subset queried by the verifier in each copy from its decision (accept or reject) on a proof that is malformed in that copy but valid in every other copy. The prover can issue such proofs for all copies (polynomially many) in parallel and learn the verifier’s query set, allowing her to break soundness after just one epoch of interaction.

We obtain the following result:

Theorem 3 (Positive result for bounded-reusable ZK-PCP, informal).

For every r , there is a ZK-PCP that is query-reusable for r epochs with query complexity $\tilde{O}(r^3)$.

Here, by query complexity, we mean the size of the query set used by the query-reusing verifier, and the $\tilde{O}$ notation hides terms that are polylogarithmic in the input length n . The theorem statement is non-trivial when the proof length n is super-cubic in r ; whereas, when $n = r^3$ for instance, the theorem is trivially satisfied by querying the whole proof. In fact, this result also holds for sublinear (ZK-)PCPs with polylogarithmic query complexity over a binary alphabet. We also show the following lower bound on the size of the query set:

Theorem 4 (Negative result for bounded-reusable ZK-PCP, informal). *Any ZK-PCP that is query-reusable for r epochs has query complexity at least $r^2/\text{polylog}(n)$.*

The above result also applies to “classical” sublinear PCPs without ZK.

From Theorems 3 and 4, there is a gap between the lower and upper bound on the query size. We leave the question of closing this gap as an interesting open problem.

1.2 Our Techniques

This section presents an overview of techniques used in our positive and negative results. We first briefly define PCP scheme and its zero-knowledge variant; refer Sect. 2.2 for formal definitions.

Let L be a language in NP. A PCP scheme for L consists of a computationally unbounded prover and a computationally bounded verifier. On input $x \in L$, prover samples a proof $\pi \xleftarrow{\$} P(x)$. The verifier samples a query set Q and a verification function represented as a circuit ver based on the input length as $(Q, \text{ver}) \xleftarrow{\$} V(1^{|x|})$. Verifier accepts the instance (x, π) if $\text{ver}(x, \pi|_Q) = 1$, where $\pi|_Q$ denotes the restriction of π to the coordinates in Q . A PCP scheme needs to satisfy the standard notions of perfect completeness and negligible soundness. In our positive results, we will use a *computationally bounded* prover with access to an NP witness w for membership of x in L .

A zero-knowledge PCP (ZK-PCP) is a PCP in which the verifier's view $(Q, \text{ver}, \pi|_Q)$ can be efficiently simulated using the input x .

The PCP is query reusable for $r(n)$ -epochs if the soundness holds for $r(n)$ -epochs when the query set and verification circuit are reused. This is formalized using the following game (defined for every n) between a malicious computationally unbounded prover A and a query reusing verifier:

1. Verifier samples $(Q, \text{ver}) \xleftarrow{\$} V(1^n)$, and reuses it throughout the game.
2. For $i = 1, \dots, r(n)$, in epoch i :
 - (a) A chooses $\{(x^{(j)}, \pi^{(j)})\}_{j \in p(n)}$ such that each $x^{(j)} \in \{0, 1\}^n$.
 - (b) A receives $b^{(j)} = \text{ver}(x^{(j)}, \pi^{(j)}|_Q)$ for each $j \in [p(n)]$.
3. A wins if it outputs (x, π) such that $x \in \{0, 1\}^n \setminus L_n$ and $\text{ver}(x, \pi|_Q) = 1$.

The PCP is query reusable for $r(n)$ -epochs if for any polynomial $p(n)$, a computationally unbounded A wins the above game with negligible probability. The above game uses only inputs of length n . This is necessary for query reusability since the query set used by the verifier depends on the input length. A formal definition of PCPs and of query reusability is given in Sect. 2.

Negative Results. Our negative results state that PCP schemes and their variants are query reusable for at most $O(\sqrt{q})$ epochs against a malicious computationally unbounded prover when the query set is of size at most q . We present an overview of our techniques proving a weaker bound on reusability; the tighter results are deferred to the technical section.

We construct an adversary A that attempts to learn a new coordinate of the query set Q reused by the verifier or substantially shrinks the set to which the reused query belongs. After learning sufficiently many coordinates of Q or sufficiently reducing the ‘entropy’ of the query set, A can obtain an instance that breaks the soundness of the scheme. We will then show that the attack can fail against an NP complete language only if a standard complexity theoretic assumption is false. This reduction is more straightforward for PCP, but more involved for the ZK-PCP. We next describe the malicious prover/attack A .

Learning Coordinates in Q Using Residual PCP. Consider the interaction between an attacker A and the verifier in the aforementioned game. At any point, the view of A consists of the instances it chose in the previous epochs and the corresponding decisions of the verifier. We define the residual PCP to be the aposteriori distribution of (Q, ver) reused by the verifier conditioned on A 's view. We will denote the residual PCP by V_{View} when View is A 's view. This name is justified by the fact that verifier outputs 1 on any instance (x, π) with probability $\Pr[\text{ver}'(x, \pi|_{Q'}) = 1 | (Q', \text{ver}') \stackrel{\$}{\leftarrow} V_{\text{View}}]$. We outline A 's strategy.

Let $\hat{Q}$ be the subset of the query set Q that A has learned in all the previous epochs. In the present epoch, A checks (i). if there exists $x \in \{0, 1\}^n \setminus L$ and a purported proof π for which the residual PCP accepts (x, π) with probability $1/3$; or (ii). if there exists $x' \in \{0, 1\}^n \cap L$ and a purported proof π' such that π' agrees with a valid proof π^* (of x') on coordinates $\hat{Q}$ but is 0 everywhere else, such that the residual PCP rejects (x', π') with probability at least $1/3$. If neither check succeeds A aborts.

If (i) succeeds, A challenges the verifier with (x, π) . If the verifier accepts (x, π) , then soundness is broken; otherwise, the set of possible queries which are consistent with A 's view—in other words, the set to which the reused query (Q, ver) belongs—shrinks by $1/3$.

Whereas, if (ii) succeeds, A challenges the verifier with (x', π') . If the verifier rejects (x', π') , then A can learn a fresh coordinate in $Q \setminus \hat{Q}$ by learning the verifier's decision on a sequence of hybrid proofs that transition from π' to π^* such that each proof differs from the previous on exactly one coordinate. There must be a pair of consecutive proofs, one accepted and the other rejected by the verifier, that disagree on exactly one coordinate. Clearly, this coordinate belongs to $Q \setminus \hat{Q}$. If (x, π) is accepted, A can cut down the set to which the reused query belongs (Q, ver) by a third.

The residual PCP should pass one of the checks (i) and (ii) unless something surprising happens: Consider an algorithm that takes an n length input and witness, samples (Q', ver') according to V_{View} , zero-pads the witness appropriately to obtain the restriction of a proof to Q' , and applies ver' to it. When the restriction of any valid PCP proof to the set $\hat{Q}$ is used as witness, this algorithm (with oracle access to the distribution V_{View}) can decide $\{0, 1\}^n \cap L$ with at most $1/3$ probability of making an error. Using a variant of Adleman's theorem, such an algorithm for an NP-complete L can be shown to contradict a standard complexity assumption, when $\hat{Q}$ is sub-linear in the input length.

The above observations imply that A breaks the soundness or learns a new coordinate in Q or substantially reduce the size of the candidate set to which the reused query belongs in every epoch. At this point, it can be argued that A breaks soundness in $O(q + n)$ epochs when the query set size q is sub-linear under a standard complexity theoretic assumption.

In Sect. 3.1, we formally present this attack with further optimizations that ensure that A breaks soundness in $q + 1$ epochs. Further, we use a more involved attack to bring down the query reusability to $\tilde{O}(\sqrt{q})$ where $\tilde{O}$ hides logarithmic factors.

For witness-indistinguishable PCP (WI-PCP) and ZK-PCP, the size of the query set need not be sub-linear in the input length. Hence, we cannot reach a contradiction using the above argument to establish that A never aborts. In fact, the negative result for WI-PCP is derived similarly to the above argument, but through a reduction to a different complexity assumption. We derive our negative result for ZK-PCP using a completely different argument which we discuss next.

Query Reusability of ZK-PCP. If A aborts with noticeable probability (i.e., checks (i) and (ii) fail), we will use the residual PCP conditioned on abort to construct an Arthur-Merlin (AM) protocol for the complement of L . This would contradict the standard complexity assumption that coNP-complete languages do not have AM protocol. Our proof technique is reminiscent of classical argument of Fortnow [11] which exploits the fact that the simulator for a statistical zero-knowledge proof behaves differently for inputs in the language and outside the language to show that the NP-complete languages do not admit statistical zero-knowledge proof.

We present the intuition behind our proof for a special case of perfect ZK-PCP schemes that additionally satisfy the following condition (that is not implied by perfect zero-knowledge): for all n and $x \in \{0, 1\}^n$, the marginal distribution of (Q, ver) according to $\text{Sim}(x)$ is identical to the distribution of the same according to $V(1^n)$. If A aborts with non-zero probability, there exists $\text{View}, \hat{Q}$ for which checks (i) and (ii) fail. Consider an AM protocol in which the prover on receiving input $x \in \{0, 1\}^n \setminus L$ chooses a private randomness s of simulator Sim for which $\text{Sim}(x; s) = (Q, \text{ver}, \pi|_Q)$ is such that (Q, ver) is consistent with View . Here, (Q, ver) is said to be consistent with View if, for each $(x', \pi, b) \in \text{View}$ (recall a view is simply the set of challenges posed by A and corresponding decision of the verifier), $\text{ver}(x', \pi|_Q) = b$. On receiving s from prover, the verifier computes $\text{Sim}(x; s) = (Q, \text{ver}, \pi|_Q)$ and checks if (Q, ver) is consistent with View . If so, verifier zero pads $\pi|_{\hat{Q}}$ (recall $\hat{Q} \subseteq Q$ for any (Q, ver) consistent with View) to obtain $\pi'|_Q$, and accepts x if $\text{ver}(x, \pi'|_Q) = 0$ and rejects otherwise.

The prover can always find s for which the associated (Q, ver) is consistent with View ; this is ensured by our assumption about the simulator. Since $x \in \{0, 1\}^n \setminus L$, for any $\pi|_Q$, $\text{ver}(x, \pi'|_Q) = 0$ with probability $2/3$, and hence verifier accepts with probability $2/3$. If $x \in \{0, 1\}^n \cap L$, if the malicious prover sends s for which the associated (Q, ver) is not consistent with View , verifier will reject. Otherwise, by perfect zero-knowledge, $\pi|_Q$ is consistent with a valid proof of x in the PCP. Hence, $\text{ver}(x, \pi'|_Q) = 0$ with probability $2/3$, and hence, verifier rejects with probability $2/3$ ensuring soundness. The verifier in this construction has oracle access to a potentially hard to sample distribution, but this can be removed using standard techniques if we settle for a non-uniform verifier, which is sufficient to contradict a standard complexity assumption.

The simple construction outlined above works for only perfect ZK-PCP with the additional unrealistic assumption. Making this intuition work for statistical ZK-PCP involves several non-trivial steps, and crucially uses a trick used in [11,

[13] based on Carter-Wegman Universal Hash Functions [5]. In the final protocol, the verifier is not silent.

Dealing with Private Randomness of Verifier. All the negative results are proved using an attack that works against a deterministic verification function ver . Indeed, the verifier's decision is deterministic in all natural PCPs for NP that we are aware of. However, in general, a query reusing verifier can use fresh private randomness when verifying every new proof while reusing the same (Q, ver) . Hence, for an instance (x, π) , the verifier's response is an independent sample from a Bernoulli distribution $\text{ver}(x, \pi|_Q, s)$ where s is sampled uniformly from a randomness domain. Our results can be extended to the general setting in a straightforward manner, although it requires several intuitive tweaks to the attack and its analysis. These tweaks are based on the simple intuition that the attacker can estimate the aforementioned Bernoulli distribution with sufficient precision using polynomially many repeated queries.

Query Reusability of Other Proof Systems. We also study reusability of input-oblivious two-message linear interactive protocol (LIP) for which query reusability can be meaningfully defined. In this proof system, the proof is a linear function π over a field $\mathbb{F}$ sampled by the prover based on the input x . The verifier's query consists of a query vector $\mathbf{v}$ and a verifier circuit ver sampled based on the input length. The verifier decision is computed as $\text{ver}(x, \pi \cdot \mathbf{v})$. Under standard complexity theoretic assumptions, we prove that, for a fixed field, LIP is not reusable even for a single epoch. We use completely different, but arguably simpler attack to show this negative result in the full-version.

Positive Results. We established that PCP and statistical ZK-PCP schemes are reusable only for $\sqrt{q}$ epochs if they query at most q coordinates of the proof. We compliment this by constructing PCP and statistical ZK-PCP schemes that are reusable for up to $q^{1/3}$ epochs ignoring logarithmic factors. Further, we achieve full query reusability for ZK-PCP by relaxing to computational setting by settling for either computational ZK-PCP or statistical ZK probabilistically checkable arguments (referred by the misnomer ZK-PCP arguments), we can achieve full reusability.

The query reusing test for zero vector testing sketched below will act as a crucial tool in all the constructions with bounded query reusability.

Tool: A Query Reusing Test for Zero Vector. Consider the property testing problem in which an n -bit string x is accepted if it is the all zero string 0^n , and rejected if it has Hamming weight more than $n/2$. The natural test for this property involves choosing $\log^2 n$ random coordinates of the given input x and accepting if and only if x is 0 in all these coordinates. The test accepts 0^n , and, by the Chernoff bound, rejects x with Hamming weight more than $n/2$ with overwhelming probability. But, the soundness of this test is broken in a single epoch since the adversary can (in parallel) learn each i in the query set from the test's output on the string that is 1 in coordinate i and 0 everywhere else.

We will first build a test that is query reusable for 1 epoch, and generalize the idea to obtain a $r(n)$ -epoch query reusable test. The test samples two subsets Q_1 and Q_2 of $[n]$ of size $\log^2 n$ uniformly at random. It accepts input x if there are no ones in $x|_{Q_1}$ and at most $\log^2 n/4$ ones in $x|_{Q_2}$ is 1, and rejects otherwise. Clearly 0^n is always accepted by the test.

We will prove that the test's output on polynomially many inputs $x_1, \dots, x_t$ chosen by the adversary in epoch 1 coincides with that of a simpler test which accepts iff there are no ones in $x|_{Q_1}$ with overwhelming probability. But then, the adversary has learned no information about Q_2 . Hence, after 1 epoch, any x such that $\|x\| \geq n/2$ chosen by the adversary after the first epoch will be rejected with overwhelming probability over the randomness of Q_2 . This implies that this test querying $2 \log^2 n$ coordinates of the input can be reused for 1 epoch.

To prove the claim, consider any arbitrary n -bit string x . The output of the test using Q_1 and Q_2 differs from that using only Q_1 only if there are no ones in $x|_{Q_1}$ and there are more than $\log^2 n/4$ ones in $x|_{Q_2}$. Since Q_1 and Q_2 are chosen uniformly and independently, the probability with which this occurs is negligible. This can be shown using Chernoff bound by separately considering inputs of weight less than $n/8$ and more than $n/8$. The claim follows by a union bound since t is polynomial in n .

A $r(n)$ -epoch query reusable test can be constructed by generalizing the above idea. The test samples $Q_1, \dots, Q_{r(n)+1}$ subsets of $[n]$ of size ℓ each (to be decided later) uniformly and independently, and accepts x if the number of ones in $x|_{Q_i}$ is less than $\ell(i-1)/2r(n)$ for all i . Using a careful analysis similar to that of the previous construction, one can show that the view of an adversary after $r(n)$ epochs can be simulated using $Q_1, \dots, Q_{r(n)}$, whenever $\ell \geq r^2(n) \log^2 n$. Thus, after $r(n)$ -epochs, the adversary does not learn any information about $Q_{r(n)+1}$, and hence fails to break the soundness of the test.

We have built a $r(n)$ -epoch query reusable zero test that queries at most $r^3(n) \log^2 n$ coordinates in the input string as long as $r(n) \leq n^{1/3-\delta}$ for some $\delta > 0$.

$r(n)$ -Epoch Query Reusable PCP ZK-PCP. Given a PCP (P, V) with query complexity $q(n)$ and proof length $\ell(n)$, we will build an $r(n)$ -epoch query reusable PCP by following the blueprint of $r(n)$ -epoch query reusable zero tester. Let $\{0, 1\}^{R_n}$ be the private randomness used by $V(n)$; i.e., $V(n)$ outputs $V(n; s)$ after sampling $s \xleftarrow{\$} \{0, 1\}^{R_n}$. The query reusing PCP uses the same prover and a new verifier V' that samples $Q_1, \dots, Q_{r(n)+1}$ where each Q_i is a random subset of $\{0, 1\}^{R_n}$ of size $\tilde{r} = r^2(n) \log^2 n$. On instance (x, π) , V' accepts if and only if, for each $i \in [r(n) + 1]$,

$$|\{s \in Q_i : (Q, \text{ver}, \text{st}) \rightarrow V(n; s), \text{ver}(x, \pi|_Q, \text{st}) = 0\}| \leq \tilde{r}(i-1)/2r(n).$$

We show a reduction from $r(n)$ -epoch query reusable zero test described above for $\{0, 1\}^{R_n}$ to $r(n)$ -epoch query reusability of (P, V') . Thus, $r(n)$ -epoch query reusability follows from that of the zero test when R_n is sufficiently large.

The $r(n)$ -epoch reusable ZK-PCP is very similar to the PCP construction. The only issue with using the same construction is that the zero-knowledge

property will be broken if the verifier queries the same proof multiple times. Hence, the prover issues $\tilde{r}(r(n) + 1)$ independent copies of the proof, so that the prover may apply each verifier to a fresh copy of the proof ensuring zero-knowledge.

Fully Query Reusable ZK-PCP. Demanding statistical zero-knowledge and soundness simultaneously for a PCP scheme leads to bounded query reusability. We realize fully query reusable computational ZK-PCP and statistical ZK-PCP arguments (where soundness is only required against bounded provers) establishing Theorem 2. Note, the soundness game for query reusability of ZK-PCP argument uses a computationally bounded adversary. Our construction uses secure Non-Interactive Secure Computation protocol with statistical semi-honest security in the OT-hybrid model (NISC/OT) [16] along with efficient implementation of non-interactive zero-knowledge proofs/arguments (NIZK) in the common reference string (CRS) model to achieve full reusability.

2 Preliminaries

2.1 Notations

We denote the set $\{1, \dots, n\}$ as $[n]$ for any $n \in \mathbb{N}$. For a vector $\mathbf{v} \in \mathbb{F}^n$, the entry in coordinate $i \in [n]$ is denoted by $\mathbf{v}_i$; similarly, for a string $x \in \{0, 1\}^n$, x_i denotes the bit in coordinate $i \in n$. Let $\pi \in \{0, 1\}^n$ and $Q \subseteq [n]$, then the restriction of π to the coordinates in Q , denoted by $\pi|_Q$ is the set $\{(i, \pi_i) : i \in Q\}$. A function $\text{negl}(n)$ is said to be negligible if, for any polynomial $p(n)$, the sequence $p(n) \cdot \text{negl}(n)$ converges to 0. Let $\mathcal{D}$ be a distribution. We use the notation $D \stackrel{\$}{\leftarrow} \mathcal{D}$ to mean that D is sampled according to the distribution $\mathcal{D}$, independent of all other random variables. For a randomized function f , $y \stackrel{\$}{\leftarrow} f(x)$ means y is assigned the randomized output of f on input x ; if f is deterministic, $y \leftarrow f(x)$ mean y is assigned the output of f on input x . When S is a set, $s \leftarrow S$ means s is sampled uniformly from S . We will succinctly describe statements of the form “probability of event A when a random variable x is sampled according to a distribution $\mathcal{D}$ conditioned on another event B is μ ” as

$$\Pr \left[x \in E \mid x \stackrel{\$}{\leftarrow} \mathcal{D}, x \in B \right] = \mu.$$

The following lemma is an application of Chernoff bound; proof is deferred to the full version.

Lemma 1. *Let $0 < \rho \leq n$, and $0 \leq \nu < \nu' \leq 1$. For any $x \in \{0, 1\}^n$,*

$$\Pr \left[(\|x|_{Q_1}\| \leq \rho\nu) \wedge (\|x|_{Q_1}\| \geq \rho\nu') \mid Q_1 \leftarrow \binom{[n]}{\rho}, Q_2 \leftarrow \binom{[n]}{\rho} \right] \leq 2e^{-\rho(\nu' - \nu)^2 / 2(\nu + \nu')}.$$

The definitions of standard complexity classes such as NP, P/poly, NP/poly and Arthur-Merlin complexity class are provided in the full version. Proposition 1 is a generalization of Adleman's theorem; refer to full version for a proof.

Proposition 1. *If $L = \{L_n\}$ is decidable using an Arthur-Merlin protocol in which the verifier has access to a distribution oracle, then L is also decidable using a non-uniform non-deterministic algorithm. In other words, L has an AM protocol in which the verifier has access to a distribution oracle if and only if $L \in \text{NP/poly}$.*

2.2 Proof Systems

Definition 1 (PCP). *Let $L = \{L_n\}_{n \in \mathbb{N}}$ be an NP language such that $L_n \subseteq \{0, 1\}^n$. A probabilistic prover P that outputs a proof π of length $\ell(n)$ on n -bit inputs, and a polynomially bounded probabilistic verifier V that outputs (Q, ver) on input 1^n , where $Q \subseteq [\ell(n)]$ is a query set of size $q(n)$ and ver is a verifier circuit, together constitute a $(\ell(n), q(n))$ -probabilistically checkable proof (PCP) scheme for L if they satisfy the following conditions:*

Completeness. For any $n \in \mathbb{N}$ and $x \in L_n$,

$$\Pr \left[\text{ver}(x, \pi|_Q) = 1 \mid (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n), \pi \stackrel{\$}{\leftarrow} P(x) \right] = 1. \quad (1)$$

Soundness. There exists a negligible function negl such that, for any $n \in \mathbb{N}$, $x \in \{0, 1\}^n \setminus L_n$ and $\pi \in \{0, 1\}^{\ell(n)}$,

$$\Pr \left[\text{ver}(x, \pi|_Q) = 1 \mid (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n) \right] \leq \text{negl}(n). \quad (2)$$

We refer to $q(n)$ and $\ell(n)$ as, respectively, query complexity and proof complexity of the PCP.

2.3 Zero Knowledge PCP

Definition 2 (ZK-PCP). *Let $L = \{L_n\}_{n \in \mathbb{N}}$ be a language in NP. (P, V) constitute a $(\ell(n), q(n))$ -zero-knowledge PCP, or simply ZK-PCP if it is a $(\ell(n), q(n))$ -PCP for L which additionally satisfies zero-knowledge property:*

Statistical Zero Knowledge. There exists a polynomial time randomized simulator Sim and a negligible function negl such that, for any n and $x \in L_n$, the total variation distance between the following distributions is at most $\text{negl}(n)$:

$$\left(Q, \pi|_Q, \text{ver} \mid (\pi|_Q, Q, \text{ver}) \stackrel{\$}{\leftarrow} \text{Sim}(x) \right), \left(Q, \pi|_Q, \text{ver} \mid \begin{array}{l} \pi \stackrel{\$}{\leftarrow} P(x) \\ (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n) \end{array} \right)$$

The ZK-PCP scheme enjoys perfect zero knowledge if negl in the above definition is the zero function.

We next define the notion of query reusability of a PCP scheme. This notion applies to WI-PCP and ZK-PCP as well.

Definition 3 (Query reusable PCP). An $(\ell(n), q(n))$ -PCP scheme (P, V) for L as defined in Definition 1 is said to be query reusable for $r(n)$ epochs if it additionally satisfies $r(n)$ -epoch soundness property:

$r(n)$ -Epoch Soundness. For any polynomial $p(n)$, there exists a negligible function $\text{negl}(n)$ such that, for any $n \in \mathbb{N}$, any computationally unbounded adversary A wins the following game with probability at most $\text{negl}(n)$:

Soundness game for r-epoch query reuse
--

- | |
|---|
| <ol style="list-style-type: none"> 1. $(Q, \text{ver}) \xleftarrow{\\$} V(1^n)$. 2. For $i = 1, \dots, r(n)$: <ol style="list-style-type: none"> (a) A chooses $\{(x^{(j)}, \pi^{(j)})\}_{j \in [p(n)]}$. (b) A receives $b^{(j)} = \text{ver}(x^{(j)}, \pi^{(j)} _Q)$ for each $j \in [p(n)]$. 3. A outputs (x^*, π^*). |
|---|

A wins the game if $x^* \in \{0, 1\}^n \setminus L_n$, and $\text{ver}(x^*, \pi^*|_Q) = 1$. We will say that a PCP scheme is fully randomness reusable if every adversary fails $r(n)$ -epoch soundness game with overwhelming probability for every polynomial $r(n)$.

Remark. To define query-reusable ZK-PCP, we must consider both $r(n)$ -epoch soundness and $r'(n)$ -epoch zero-knowledge. However, $r'(n)$ -epoch zero-knowledge is only needed for our positive results. For negative results, $r(n)$ -epoch soundness alone suffices: if it fails, the ZK-PCP is not query-reusable. Using this weaker definition strengthens our negative results. Thus, we defer the definition of $r'(n)$ -epoch zero-knowledge to Sect. 4.3 (see Definition 5).

3 Negative Results

Consider an adversary that interacts with a query reusing verifier. At any stage of the interaction, the view of the adversary consists of the instances it chose along with the verifier's output on those instances. We refer to the aposteriori distribution of the query that is being reused by the verifier as the residual PCP, since the distribution of the future outputs of the query reusing verifier is completely decided by this distribution. The description of the prover is not included in the description of the residual PCP since it remains the same as that of the parent PCP. We formally define the notion below.

Definition 4. Let (P, V) be a $(\ell(n), q(n))$ -PCP and $m, n \in \mathbb{N}$. Consider an adversary who learns a query reusing verifier's output $b^{(i)} \in \{0, 1\}$ on instance $x^{(i)}, \pi^{(i)}$ where $x^{(i)} \in \{0, 1\}^n$ for each $i \in [m]$. We denote the view of such an adversary as $\text{View} = \{(x^{(i)}, \pi^{(i)}, b^{(i)}) : i \in [m]\}$. The residual PCP scheme for n -bit inputs, conditioned on View , denoted by V_{View} : we outputs (Q, ver) with

probability

$$\Pr \left[(Q, \text{ver}) = (Q', \text{ver}') \mid \begin{array}{c} (Q', \text{ver}') \xleftarrow{\$} V(1^n), \\ \text{ver}'(x^{(i)}, \pi^{(i)}|_{Q'}) = b^{(i)} \forall i \in [m] \end{array} \right] \quad (3)$$

Proposition 2. Let $\text{View} = \{x^{(i)}, \pi^{(i)}, b^{(i)}\}_{i \in [m]}$, where $x^{(i)} \in \{0, 1\}^n$ for all $i \in [m]$. Then, for any $x \in \{0, 1\}^n$, and $\pi \in \{0, 1\}^{\ell(n)}$,

$$\begin{aligned} \Pr \left[\text{ver}(x, \pi|_Q) = 1 \mid (Q, \text{ver}) \xleftarrow{\$} V(1^n), \text{ver}(x^{(i)}, \pi^{(i)}|_Q) = b^{(i)} \forall i \in [m] \right] \\ = \Pr \left[\text{ver}(x, \pi|_Q) = 1 \mid (Q, \text{ver}) \xleftarrow{\$} V_{\text{View}} \right]. \end{aligned}$$

The above proposition justifies the definition of residual PCP since it establishes that the correctness and soundness of the query reusing PCP conditioned on adversary's view is the same as that of the residual PCP scheme. The residual PCP scheme inherits the completeness property of its parent PCP scheme since the completeness property always holds. However, the soundness error of a residual tester may be arbitrarily large.

3.1 Attack Against Query Reusable PCP

Lemma 2. Let $\tilde{V}$ be a residual PCP scheme for n -length inputs corresponding to a $(\ell(n), q(n))$ -PCP scheme V . For any $\hat{Q} \subseteq [\ell(n)]$, there exists an algorithm that outputs $x^{(i)} \in \{0, 1\}^n$ and $\pi^{(i)} \in \{0, 1\}^\ell$ for each $i \in [t]$, where $t \leq 2n^2$, and functions f, g, h such that one of the following events occurs with at least $1 - (2/3)^{-n}$ probability over the randomness of $(Q, \text{ver}) \xleftarrow{\$} \tilde{V}$ when $b^{(i)} \leftarrow \text{ver}(x^{(i)}, \pi^{(i)}|_Q)$ for all $i \in [t]$.

1. $f(\{(x^{(i)}, \pi^{(i)}, b^{(i)})\}_{i \in [t]}) \in Q \setminus \hat{Q}$.
2. $x \in \{0, 1\}^n \setminus L_n$ and $\text{ver}(x, \pi|_Q) = 1$ when $g(\{(x^{(i)}, \pi^{(i)}, b^{(i)})\}_{i \in [t]}) = (x, \pi)$.
3. If $h(\{(x^{(i)}, \pi^{(i)}, b^{(i)})\}_{i \in [t]}) = 1$, for any $x \in L_n$ and any valid proof π ,

$$\Pr \left[\text{ver}'(x, (\pi|_{\hat{Q}}, 0^n|_{[\ell] \setminus \hat{Q}})|_{Q'}) = 1 \mid \begin{array}{c} (Q', \text{ver}') \xleftarrow{\$} \tilde{V} \\ \text{ver}'(x^{(i)}, \pi^{(i)}|_{Q'}) = b^{(i)} \forall i \in [t] \end{array} \right] \geq 2/3,$$

and for all $x \in \{0, 1\}^n \setminus L_n$ and π ,

$$\Pr \left[\text{ver}'(x, \pi|_{Q'}) = 1 \mid \begin{array}{c} (Q', \text{ver}') \xleftarrow{\$} \tilde{V} \\ \text{ver}'(x^{(i)}, \pi^{(i)}|_{Q'}) = b^{(i)} \forall i \in [t] \end{array} \right] \leq 1/3,$$

Additionally, f and g guarantee that if their outputs are not $\perp$, then necessarily the first and second conditions, respectively, are met.

Proof. The algorithm works as follows:

Description of Algorithm in Lemma 2

1. Set $S \leftarrow \emptyset, C \leftarrow \emptyset$ and $k \leftarrow 1$.
2. Set $j \leftarrow 1$.
3. If there exists $\hat{x} \in \{0, 1\}^n \setminus L_n, \hat{\pi} \in \{0, 1\}^\ell$ such that

$$\Pr \left[\text{ver}(\hat{x}, \hat{\pi}|_Q) = 1 \left| \begin{array}{l} (Q, \text{ver}) \stackrel{\$}{\leftarrow} \tilde{V} \\ \text{ver}(x, \pi|_Q) = 0 \forall (x, \pi) \in S \\ \text{ver}(x, \pi|_Q) = 1 \forall (x, \pi) \in C \end{array} \right. \right] \geq 1/3, \quad (4)$$

(when many exist choose one deterministically) then set $S \leftarrow S \cup \{(\hat{x}, \hat{\pi})\}$ and increment j . If $j \leq n$ got to the beginning of this step; else if $j = n + 1$ and $k \leq n$, go to next step; finally, if $j = n + 1$ and $k = n + 1$, go to step 5. Otherwise, if Equation (4) does not hold for any $x \in \{0, 1\}^n \setminus L_n$ and $\pi \in \{0, 1\}^\ell$ go to the next step.

4. If there exists $\hat{x} \in L_n$, and a valid proof $\hat{\pi}$ such that

$$\Pr \left[\text{ver}(\hat{x}, (\hat{\pi}|_{\hat{Q}}, 0^n|_{[\ell] \setminus \hat{Q}})|_Q) = 1 \left| \begin{array}{l} (Q, \text{ver}) \stackrel{\$}{\leftarrow} \tilde{V} \\ \text{ver}(x, \pi|_Q) = 0 \forall (x, \pi) \in S \\ \text{ver}(x, \pi|_Q) = 1 \forall (x, \pi) \in C \end{array} \right. \right] \leq 2/3, \quad (5)$$

(when many such $(\hat{x}, \hat{\pi})$ exist choose one deterministically) then set

$$C \leftarrow C \cup \{(\hat{x}, \pi^{(i)}) : i \in [\ell - |\hat{Q}|], \pi^{(i)} = (\hat{\pi}|_{[\ell] \setminus \hat{Q}_i}, 0^n|_{\hat{Q}_i})\}$$

where, $\hat{Q}_i = \{q_1, \dots, q_i\}$ when $[\ell] \setminus \hat{Q} = \{q_1, \dots, q_{\ell - |\hat{Q}|}\}$. Increment k and go to step 2. Otherwise, if Equation (5) does not hold for any $x \in L_n$ and any of its valid proof π , go to the next step.

5. Let $t = |C| + |S| \leq n^2 + n\ell(n)$. Output $\{(x^{(i)}, \pi^{(i)})\}$ such that $(x^{(i)}, \pi^{(i)})$ is a distinct member of C for each $i \leq |C|$ and a distinct member of S for each $|C| < i \leq t$.

For each $i \in [t]$, let $b^{(i)} \leftarrow \text{ver}(x^{(i)}, \pi^{(i)}|_Q)$ when $(Q, \text{ver}) \stackrel{\$}{\leftarrow} \tilde{V}$.

Suppose there exists $(x^{(i)}, \pi^{(i)}) \in C$ such that $b^{(i)} = 0$. Then, there exists $x^* \in L_n$, a valid proof π^* , and $j \in [\ell - |\hat{Q}|]$ such that

$$\begin{aligned} \text{ver}(x^*, (\pi^*|_{[\ell] \setminus \{q_1, \dots, q_{j-1}\}}, 0^n|_{\{q_1, \dots, q_{j-1}\}})|_Q) &= 1 \\ \text{and } \text{ver}((\pi^*|_{[\ell] \setminus \{q_1, \dots, q_j\}}, 0^n|_{\{q_1, \dots, q_j\}})|_Q) &= 0. \end{aligned}$$

To establish the above statement for the case where $j = 1$, we used the fact that, since π^* is a valid proof, $\text{ver}(x^*, \pi^*|_Q) = 1$. But then, $q_j \in Q \setminus \hat{Q}$. We set $f(\{(x^{(i)}, \pi^{(i)}, b^{(i)})\})$ to q_j in this case, and to $\perp$ otherwise. Thus, whenever the output of f is not $\perp$, the first condition in the lemma is satisfied.

Set $g(\{(x^{(i)}, \pi^{(i)}, b^{(i)})\}) = (x^{(i)}, \pi^{(i)})$ if there exists $i > |C|$ such that $b^{(i)} = 1$, and $\perp$ otherwise. Thus, whenever the output of g is $(x^{(i)}, \pi^{(i)})$ (and not $\perp$), we have $x^{(i)} \in \{0, 1\}^n \setminus L_n$ and $\text{ver}(x^{(i)}, \pi^{(i)}|_Q) = 1$ satisfying condition 2 in the lemma.

Finally, set $h(\{(x^{(i)}, \pi^{(i)}, b^{(i)})\}) = (x^{(i)}, \pi^{(i)}) = 1$ if the algorithm terminates with values $j \leq n$ and $k \leq n$, and $b^{(i)} = 1$ for all $i \leq |C|$ and $b^{(i)} = 0$ for all $i > |C|$. Set the output of h to be $\perp$ otherwise. This implies that the value of C and S at the time of termination of the algorithm are such that (see step 3) Eq. (4) does not hold for any $x \in \{0, 1\}^n \setminus L_n$ and $\pi \in \{0, 1\}^\ell$, and (see step 4) Eq. (5) does not hold for any $x \in L_n$ and any of its valid proof π . This directly implies condition 3 in the lemma.

The above three definitions imply that, if none of the three conditions in the lemma are met, then the algorithm terminates with values $j = n+1$ or $k = n+1$, and $b^{(i)} = 1$ for all $i \leq |C|$ and $b^{(i)} = 0$ for all $i > |C|$. The theorem follows from the following claim:

Claim. Suppose the algorithm terminates with values $j = n+1$ or $k = n+1$. When $(Q, \text{ver}) \stackrel{\$}{\leftarrow} \tilde{V}$, and $b^{(i)} \leftarrow \text{ver}(x^{(i)}, \pi^{(i)}|_Q)$ for each $i \in [t]$,

$$\Pr \left[(b^{(i)} = 1, \forall i \leq |C|) \wedge (b^{(i)} = 0, \forall i > |C|) \right] \leq (2/3)^{-n}.$$

The claim is proved by making repeated application of Eq. (4) and Eq. (5). We provide a formal proof in the full version.

For each n , consider the following deterministic adversary A that interacts with an $r(n)$ -epoch reusable PCP scheme V as follows in the $r(n)$ -epoch soundness game defined in Definition 3. Having fixed n , we will suppress n when there is no confusion.

Description of Adversary A against Reusable PCP

1. $(Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n)$.
2. Set $\text{View} \leftarrow \emptyset, \hat{Q} \leftarrow \emptyset, \text{ep} \leftarrow 1$.
3. Invoke Lemma 2 with $V = V_{\text{View}}$ and $\hat{Q}$ to obtain $\{(x^{(i)}, \pi^{(i)})\}_{i \in [t]}$ for $t \leq 2n^2$. A uses $\{(x^{(i)}, \pi^{(i)})\}_{i \in [t]}$ as input to V in epoch ep .
4. For each $i \in [t]$, A receives $b^{(i)} \leftarrow \text{ver}(x^{(i)}, \pi^{(i)}|_Q)$ and sets $\text{View} \leftarrow \text{View} \cup \{(x^{(i)}, \pi^{(i)}, b^{(i)})\}_{i \in [t]}$. Compute the output of the functions f, g, h defined in Lemma 2 using $\tilde{V}$ and $\{(x^{(i)}, \pi^{(i)}, b^{(i)})\}_{i \in [t]}$.
 - (a) If g outputs (x, π) (and not $\perp$), then A outputs (x, π) and terminates.
 - (b) Else if g outputs $\perp$ and f outputs $j \in [\ell]$ (and not $\perp$), set $\hat{Q} \leftarrow \hat{Q} \cup \{j\}$.
 - (c) Else if f and g output $\perp$, and h outputs 1, then A outputs $\perp$ and terminates.
 - (d) Else if f, g and h output $\perp$, then A outputs $\perp$ and terminates.
5. If $\text{ep} = r$, output $\perp$ and terminate; else increment ep , and go to step 3.

We will use the following notation in the sequel. Let View be a view, and (Q, ver) be a query in the support of $V(1^n)$. We will say that (Q, ver) is consistent with

View if, for every $(x, \pi, b) \in \text{View}$, $\text{ver}(x, \pi|_Q) = b$. We will abuse denote this relationship as $(Q, \text{ver}) \sim \text{View}$.

Lemma 3. *Let (P, V) be a $(\ell(n), q(n))$ -PCP system for a language $L = \{L_n\}$. For any $n \in \mathbb{N}$, when $r(n) = q(n) + 1$, suppose adversary A fails the $r(n)$ -epoch soundness game for L_n against query reusing verifier with noticeable probability $\mu(n)$. Then, there exists $m(n) \geq 1$ and $\{\text{View}_i\}_{i \in [m(n)]}$ such that:*

1. *Any (Q, ver) in the support of $V(1^n)$ is consistent with at most one View_i . Further,*

$$\Pr[\exists i \in [m], (Q, \text{ver}) \sim \text{View}_i | (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n)] \geq \mu(n) - r(n)(2/3)^{-n}.$$

2. *For each $i \in [m]$, there exists $\hat{Q} \subset [\ell(n)]$ such that, for all $(Q, \text{ver}) \sim \text{View}_i$, $\hat{Q} \subseteq Q$, and*
 - (a) *For all $x \in \{0, 1\}^n \setminus L$ and $\pi \in \{0, 1\}^\ell$,*

$$\Pr[\text{ver}(x, \pi|_Q) = 1 | (Q, \text{ver}) \stackrel{\$}{\leftarrow} V_{\text{View}_i}] \leq 1/3. \quad (6)$$

- (b) *For any $x \in L_n$ and any valid proof π ,*

$$\Pr[\text{ver}(x, (\pi|_{\hat{Q}}, 0^n|_{[\ell] \setminus \hat{Q}})|_Q) = 1 | (Q, \text{ver}) \stackrel{\$}{\leftarrow} V_{\text{View}_i}] \geq 2/3. \quad (7)$$

Proof. Fix $n \in \mathbb{N}$. We will suppress n in the notations when there is no confusion. For any value of $\text{ep} \leq r$, whenever A outputs (x, π) in step 4(a), $x \in \{0, 1\}^n \setminus L_n$ and $\text{ver}(x, \pi|_Q) = 1$. This follows from statement 2 of Lemma 2. Hence, such (x, π) wins the soundness game. Thus, A fails only if it outputs $\perp$.

We will next establish that the instruction in step 5 to output $\perp$ and terminate if $\text{ep} = r$ is never executed. It suffices to show that, in step 4 of epoch r ($\text{ep} = r$), A will output (x, π) in step 4.(a) or terminate in step 4.(c) or 4.(d); i.e., A does not enter step 4.(b). To see this, observe that, in each epoch $\text{ep} < r$, A either terminates in 4.(a), 4.(c) or 4.(d), or adds a new element to $\hat{Q}$ in step 4.(b). By condition 2 of Lemma 2, every new element added to $\hat{Q}$ belongs to $Q \setminus \hat{Q}$, where Q is as defined in step 1. Thus, when A enters epoch r , $\hat{Q} = Q$ as $|Q| = r - 1$. The observation now follows. Along with the previous observation, this implies that A fails if and only if it terminates in step 4.(c) or 4.(d) in some epoch.

Next, we will argue that, over the randomness of (Q, ver) chosen by $V(1^n)$, A terminates in step 4.(d) with negligible probability. Since the algorithm in Lemma 2 and A are deterministic, the values of all variables at any point in the execution of A is a deterministic function of (Q, ver) . Thus, specifically, the value of $\tilde{V}, \hat{Q}$ in any epoch $1 \leq \text{ep} \leq r$ of the execution of A is unique for every (Q, ver) . Fix any epoch $1 \leq \text{ep} \leq r$. Let $\mathcal{T}$ be the set of all possible values that

the pair $(\tilde{V}, \hat{Q})$ can take in epoch ep . Define E as the event “ A terminates in step 4.(c) in epoch ep ”. Then, over the randomness of $(Q, \text{ver}) \xleftarrow{\$} V(1^n)$,

$$\Pr[E] = \sum_{(a,b) \in \mathcal{T}} \Pr[(\tilde{V}, \hat{Q}) = (a, b)] \cdot \Pr[E | (\tilde{V}, \hat{Q}) = (a, b)].$$

When $(\tilde{V}, \hat{Q}) = (a, b)$, the event E occurs only if the output of algorithm in Lemma 2 for $(\tilde{V}, \hat{Q}) = (a, b)$ does not satisfy any of the three conditions in the lemma. This occurs with probability at most $(2/3)^{-n}$ by the same lemma. Hence, by the above equation, the probability with which (Q, ver) chosen by $V(1^n)$ causes A to terminate in step 4.(c) in epoch ep is $(2/3)^{-n}$, for any $1 \leq \text{ep} \leq r$. By a union bound, A terminates in step 4.(c) (in any epoch) with at most $r(2/3)^{-n}$.

Let $\mathcal{V}$ be the set of all possible final view of A . Define

$$\mathcal{G} = \{\text{View} \in \mathcal{V} : A \text{ terminates in step 4.(c) with View as final view}\}.$$

Let $\mathcal{G} = \{\text{View}_1, \dots, \text{View}_m\}$ for some m . Each View_i is a final view of A ; in that, A terminates when its view is View_i . Since the view of A at any point of the attack is a deterministic function of (Q, ver) chosen by the verifier, the final view of A is a deterministic function of (Q, ver) chosen by the verifier. Thus, $(Q, \text{ver}) \sim \text{View}_i$ for at most one such i .

For any $i \in [m]$, we will associate View_i with the set $\hat{Q}$ in the final state of A when it terminates in step 4.(c). Then, if $(Q, \text{ver}) \sim \text{View}_i$, then $\hat{Q} \subseteq Q$. This is guaranteed by the property of f stated in condition 1 of Lemma 2. Finally, A terminates in step 4.(c) when h outputs 1. Hence, statements 2.(a) and 2.(b) in the lemma follows from statement 3 of Lemma 2, and the value of variables $\tilde{V}, \hat{Q}$ and View at the time of termination. Thus, the statement 2 of the lemma holds.

3.2 Query Reusability of PCP

Theorem 5. *Let (P, V) be a $(\ell(n), q(n))$ -PCP system for a language L . If L does not belong to NP/poly with $q(n)$ -bit long witness, then there exists an adversary that wins the $q(n) + 1$ -epoch soundness game with overwhelming probability for by making $O(n\ell(n))$ evaluations in each epoch.*

Proof. Fix $n \in \mathbb{N}$. Consider the $q(n) + 1$ -epoch soundness game between A and $V(1^n)$. If A fails with noticeable probability, by Lemma 3, there exists View_1 and $\hat{Q}$ that satisfy the second condition in the lemma. We will build a probabilistic polynomial time algorithm B with access to a distribution oracle that takes $|\hat{Q}| \leq q(n)$ sized witness, and correctly decides L_n with probability more than $2/3$.

For any $x \in L_n$, define the witness to be $\pi|_{\hat{Q}}$, for any valid proof π . On receiving $x \in \{0, 1\}$ and witness of the form $\pi|_{\hat{Q}}$, B then the distribution oracle to sample $(Q, \text{ver}) \xleftarrow{\$} V(1^n)$ subject to $(Q, \text{ver}) \sim \text{View}_1$. B outputs

$\text{ver}(x, (\pi|_{\hat{Q}}, 0^n|_{[\ell] \setminus \hat{Q}})|_Q)$. By the second statement in Lemma 3, the output is 1 with probability at least $2/3$ if $x \in L_n$ and $\pi|_{\hat{Q}}$ is a valid witness; whereas, for any $x \in \{0, 1\}^n \setminus L_n$ and any purported witness, output is 1 with probability at most $1/3$.

Using Adleman's theorem, we can show that B implies a polynomial time algorithm that takes $|\hat{Q}|$ sized witness to decide L_n .

Thus, if L does not belong to NP/poly with $q(n)$ -bit long witness, A cannot fail the $q(n)+1$ -epoch soundness game with noticeable probability. The theorem follows from the fact that A makes $O(n\ell(n))$ evaluations in each epoch.

Using a similar analysis, we can obtain a negative result for witness indistinguishable PCP. Specifically, we can show that perfect WI-PCP with query complexity $q(n)$ is not reusable for $q(n)$ epochs unless non-interactive statistical witness indistinguishability exists for NP. We formally define WI-PCP and prove this result in the full version.

3.3 Query Reusability of ZK-PCP

Theorem 6. *Let (P, V) be a statistical $(\ell(n), q(n))$ -ZK-PCP system for an NP-complete language L . If coNP is not contained in NP/poly, there exists an adversary that wins the $q(n)+1$ -epoch soundness game with overwhelming probability by making $O(n\ell(n))$ evaluations in each epoch.*

Proof. Let L be NP-complete. Let Sim be the statistical simulator for (P, V) with zero-knowledge error at most $\text{negl}(n)$, where negl is negligible. Suppose the attack A described in Sect. 3.1 fails the $q(n)+1$ -epoch soundness game for n -bit inputs with noticeable probability $\mu(n)$. We will build a AM protocol that decides $\{0, 1\}^n \setminus L$ in which verifier has oracle access to a (hard to compute) distribution. But then, by Proposition 1, the complement of L belongs to NP/poly contradicting the standard assumption that coNP is not contained in NP/poly.

Fix n . When A fails with probability μ , there exist $\text{View}_1, \dots, \text{View}_m$ that satisfy the properties stated in Lemma 3. For each $i \in [m]$, let $\hat{Q}_i$ be and set $\hat{Q}$ associated with View_i in Lemma 3. Let $\mathcal{I}$ be the distribution that chooses $i \in [m]$ with probability

$$\Pr[(Q, \text{ver}) \sim \text{View}_i | (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n), \exists j \in [m] : (Q, \text{ver}) \sim \text{View}_j].$$

Let $\mathcal{D}$ be a distribution as follows with respect to a function $f : [m] \rightarrow \mathbb{N} \times \mathbb{N}$ that will be define at the end of the proof.

$$\left(i, Q, \text{ver}, \hat{Q}_i, l_i, b_i \left| \begin{array}{l} i \stackrel{\$}{\leftarrow} \mathcal{I}, (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n), \\ (Q, \text{ver}) \sim \text{View}_i, (l_i, b_i) \leftarrow f(i) \end{array} \right. \right) \quad (8)$$

Consider the following AM protocol (P', V') in which V' is given oracle access to distribution $\mathcal{D}$:

1. On input $x \in \{0, 1\}^n \setminus L$, P' checks if the distribution of (Q, ver) according to $\text{Sim}(x)$ is more than $(\mu/4)^2$ far in total variation distance from that according to $V(1^n)$. If so, go to step 2; else go to step 3.
 2. P' invites V' to participate in an AM sub-protocol which accepts a pair of circuits if their output distributions are more than $(\mu/4)^2$ apart in total variation distance with probability $2/3$, and rejects if they are at most $\text{negl}(n)$ apart with probability $2/3$. P' and V' engage in the protocol with the circuits sampling the outputs of $\text{Sim}(x)$ (with x hard coded) and $V(1^n)$ as the common input. V' accepts x , if the AM protocol accepts the circuits, and rejects otherwise.
 3. In the following steps, P' invites V' to participate in an AM sub-protocol which accepts $x \in \{0, 1\}^n \setminus L_n$ for which the distribution of (Q, ver) according to $\text{Sim}(x)$ and according to $V(1^n)$ are at most $(\mu/4)^2$ apart in total variation distance with probability at least $5/9$, and accepts any $x \in L_n$ with probability at most $4/9$.
 4. V' invokes $\mathcal{D}$ to receive $(i, Q, \text{ver}, \hat{Q}_i, l_i, b_i)$. V' samples l_i independent random linear transformations $f_1, \dots, f_{l_i} : \{0, 1\}^R \rightarrow \{0, 1\}^{b_i}$, and l_i^2 points $z_1, \dots, z_{l_i^2} \in \{0, 1\}^{b_i}$. V' sends i and $\{f_j\}_{j \in [l_i]}$ and $\{z_j\}_{j \in [l_i^2]}$ to P' .
 5. Let $\{0, 1\}^R$ be the randomness space of the simulator Sim for input length n . P' sends $s \in \{0, 1\}^R$ such that $f_i(r) = z_j$ for some $i \in [l_i]$ and $j \in [l_i^2]$, and $(Q', \text{ver}') \sim \text{View}_i$ where $(Q', \text{ver}', \pi|_{Q'}) \leftarrow \text{Sim}(x; s)$. If no such s exists, V' sends $r = \perp$.
 6. If $r = \perp$ or there exists no $j \in [l_i]$ and $k \in [l_i^2]$ such that $f_j(r) = z_k$ or $(Q', \text{ver}') \not\sim \text{View}_i$ when $(Q', \text{ver}', \pi|_{Q'}) \leftarrow \text{Sim}(x; s)$, then V' rejects x . Here, checking if $(Q', \text{ver}') \sim \text{View}_i$ can be carried out efficiently since this amounts to checking the output of ver' on polynomially many instances.
 7. Otherwise, V' accepts if $\text{ver}(x, (\pi|_{\hat{Q}_i}, 0^n|_{Q \setminus \hat{Q}_i})|_Q) = 0$, and rejects otherwise.
- Recall, $\pi|_{\hat{Q}_i}$ is well defined since $\hat{Q}_i \subseteq Q'$ as $(Q', \text{ver}') \sim \text{View}_i$.

Lemma 4. *If $x \in \{0, 1\}^n \setminus L$ such that the distribution of (Q, ver) according to $\text{Sim}(x)$ is more than $(\mu/4)^2$ far in total variation distance from that according to $V(1^n)$, then the above protocol accepts x with probability $2/3$. Further, if $x \in L_n$ and a malicious prover invites V' to execute step 2, then V' rejects with probability $2/3$.*

Proof. The proof follows from the guarantee of the AM sub-protocol and the fact that, when $x \in L_n$, the distribution of (Q, ver) according to $\text{Sim}(x)$ and that according to $V(1^n)$ are at most $\text{negl}(n)$ apart in total variation distance.

Lemma 5. *In the AM protocol described in steps 3–7, V' accepts $x \in \{0, 1\}^n \setminus L_n$ for which the distribution of (Q, ver) according to $\text{Sim}(x)$ and that according to $V(1^n)$ are at most $(\mu/4)^2$ apart in total variation distance with probability at least $5/9$, and accepts any $x \in L_n$ with probability at most $4/9$.*

The theorem follows from Lemma 4 and Lemma 5.

Proof (Proof of Lemma 5). Suppose $x \in \{0, 1\}^n \setminus L$. If P' sends $s \in \{0, 1\}^R$ that succeeds the checks in step 6, then V' will accept with $2/3$ probability. This follows from Lemma 3 which guarantees Eq. (6) for any $i \in [m]$ and $\pi \in \{0, 1\}^{\ell(n)}$. We will show that, when the distribution of (Q, ver) according to $\text{Sim}(x)$ and that according to $V(1^n)$ are at most $(\mu/4)^2$ apart in total variation distance, over the randomness of $\{f_i\}$ and $\{z_i\}$ sent by V' (in step 4), s sent by P' succeeds the checks with substantial probability.

Let $\text{Sup}(P(x))$ denote the set of all possible proofs that can be sampled (with non-zero probability) by the honest prover P on input $x \in L_n$. When $x \in L_n$, V' will (i) reject if the malicious prover sends $s \in \{0, 1\}^R$ that fails the checks in step 6, and (ii) reject with probability $2/3$ if s succeeds the checks in step 6, but there exists $\pi' \in \text{Sup}(P(x))$ such that $\pi'|_{\hat{Q}_i} = \pi|_{\hat{Q}_i}$ when $(Q, \text{ver}, \pi|_Q) \leftarrow \text{Sim}(x; s)$. Here, (ii) follows from Lemma 3 which guarantees Eq. (7) for any $i \in [m]$ and a valid proof π' . We crucially used the fact that, when $\pi' \in \text{Sup}(P(x))$, π' is indeed such a proof. We will show that, with substantial probability over the randomness of $\{f_i\}$ and $\{z_i\}$ sent by V' (in step 4), there exists no $s \in \{0, 1\}^R$ that does not meet the conditions in (i) or (ii).

We will prove the above two statements, which will prove the lemma. The following definition will be used in the sequel:

$$\nu_i = \Pr \left[(Q, \text{ver}) \sim \text{View}_i \mid (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n) \right] \quad \forall i \in [m]. \quad (9)$$

Claim. Fix any $x \in \{0, 1\}^n \setminus L$ for which the distribution of (Q, ver) according to $\text{Sim}(x)$ and that according to $V(1^n)$ are at most $(\mu/4)^2$ apart in statistical distance. Define $\mathcal{G}_x$ as the set of all $i \in [m]$ such that

$$\Pr \left[(Q, \text{ver}) \sim \text{View}_i \mid s \stackrel{\$}{\leftarrow} \{0, 1\}^R, (Q, \text{ver}, \pi|_Q) \leftarrow \text{Sim}(x; s) \right] \geq \nu_i/2.$$

Then, $\Pr[i \in \text{Good}_x \mid i \stackrel{\$}{\leftarrow} \mathcal{I}] \geq 1 - (\mu/4)$.

Proof. Since $\Pr[(Q', \text{ver}') \sim \text{View}_i \mid (Q', \text{ver}') \stackrel{\$}{\leftarrow} V(1^n)] = \nu_i$ for all $i \in [m]$, by the assumption in the claim for $x \in \{0, 1\}^n \setminus L$,

$$\begin{aligned} & \frac{1}{2} \cdot \Pr \left[\exists i \in [m] \setminus \mathcal{G}_x \text{ s.t. } (Q, \text{ver}) \sim \text{View}_i \mid (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n) \right] \\ & \leq \Pr \left[\exists i \in [m] \setminus \mathcal{G}_x \text{ s.t. } (Q, \text{ver}) \sim \text{View}_i \mid (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n) \right] \\ & \quad - \Pr \left[\exists i \in [m] \setminus \mathcal{G}_x \text{ s.t. } (Q, \text{ver}) \sim \text{View}_i \mid \begin{array}{c} s \stackrel{\$}{\leftarrow} \{0, 1\}^R, \\ (Q, \text{ver}, \pi|_Q) \leftarrow \text{Sim}(x; s) \end{array} \right] \\ & \leq (\mu/4)^2. \end{aligned}$$

The final inequality used the definition of total variation distance. By Lemma 3, (Q, ver) sampled according to $V(1^n)$ is consistent with View_i for some $i \in [m]$ with probability at least $(\mu/2)$. Hence, by the above inequality,

$$\Pr \left[\exists i \in [m] \setminus \mathcal{G}_x \text{ s.t. } (Q, \text{ver}) \sim \text{View}_i \mid \begin{array}{c} (Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n) \\ \exists i \in [m] \text{ s.t. } (Q, \text{ver}) \sim \text{View}_i \end{array} \right] \leq (\mu/4). \quad (10)$$

This proves the claim.

Claim. Fix any $x \in L_n$. Define $\mathcal{B}_x$ as the set of all $i \in [m]$ such that,

$$\Pr \left[\begin{array}{c} (Q, \text{ver}) \sim \text{View}_i \\ \pi|_{\hat{Q}_i} \neq \pi'|_{\hat{Q}_i} \forall \pi' \in \text{Sup}(P(x)) \end{array} \middle| \begin{array}{c} s \xleftarrow{\$} \{0, 1\}^R, \\ (Q, \text{ver}, \pi|_Q) \leftarrow \text{Sim}(x; s) \end{array} \right] \geq \nu_i / 16^9 R^3.$$

Then, $\Pr[i \in \text{Bad}_x | i \xleftarrow{\$} \mathcal{I}] \leq \text{negl}'(n)$ for a negligible function negl' .

The proof uses a similar argument as in the proof of the previous claim. We defer the proof to the full version.

At this point, we will use the following known result about Carter-Wegman Universal Hash Functions, taken from [13].

Lemma 6. *Let $f_1, \dots, f_l : \{0, 1\}^R \rightarrow \{0, 1\}^b$ be l uniform independent linear transformations, and $z_1, \dots, z_{l^2} \in \{0, 1\}^b$ be l^2 uniform independent vectors. If $l > \max(16, b)$,*

1. *For any $\mathcal{T} \subseteq \{0, 1\}^R$ such that $4|\mathcal{T}| \leq 2^b \leq 16|\mathcal{T}|$,*

$$\Pr[\exists s \in \mathcal{T}, i \in [l], j \in [l^2] : f_i(s) = z_j] \geq 1 - 2^{-l/16}.$$

2. *For any $\mathcal{T}' \subseteq \{0, 1\}^R$ such that $|\mathcal{T}'| \leq 2^b/d$,*

$$\Pr[\exists s \in \mathcal{T}', i \in [l], j \in [l^2] : f_i(s) = z_j] \leq l^3/d.$$

Define f such that $f(i) = (l_i, b_i)$ such that $4\nu_i \cdot 2^R \leq 2^{b_i} \leq 8\nu_i \cdot 2^R$, and $l_i = \max(16^2, b_i)$ for all $i \in [m]$. For any $x \in \{0, 1\}^n \setminus L_n$ and $i \in [m]$, define

$$\mathcal{T}_{x,i} = \{s \in \{0, 1\}^R : (Q, \text{ver}, \pi|_Q) \leftarrow \text{Sim}(x; s), (Q, \text{ver}) \sim \text{View}_i\}.$$

When $i \in \mathcal{G}_x$, $4|\mathcal{T}_{x,i}| \leq 2^{b_i} \leq 16|\mathcal{T}_{x,i}|$. Next, for any $x \in L_n$ and $i \in [m]$, define

$$\mathcal{T}'_{x,i} = \left\{ s \in \{0, 1\}^R : \begin{array}{c} (Q, \text{ver}, \pi|_Q) \leftarrow \text{Sim}(x; s), \\ (Q, \text{ver}) \sim \text{View}_i, \\ \pi|_{\hat{Q}_i} \neq \pi'|_{\hat{Q}_i} \forall \pi' \in \text{Sup}(P(x)) \end{array} \right\}.$$

When $i \notin \mathcal{B}_x$, $|\mathcal{T}'_{x,i}| \leq \frac{\nu_i \cdot 2^R}{16^9 R^3} \leq \frac{2^{b_i}}{16^3 R^3} \leq \frac{2^{b_i}}{16^3 \cdot l_i^3}$.

For any $x \in \{0, 1\}^n \setminus L$ and $i \in \mathcal{G}_x$, appealing to the first statement in Lemma 6 invoked with $l = l_i, b = b_i$ and $\mathcal{T} = \mathcal{T}_{x,i}$, we conclude that, when V' sends i in step 4, over the randomness of $\{f_j\}, \{z_j\}$, an honest prover can find $s \in \{0, 1\}^R$ that succeeds the checks in step 6 with probability at least $1 - 2^{-16}$. On the other hand, for any $x \in L_n$ and $i \in [m] \setminus \mathcal{B}_x$, appealing to the second statement in Lemma 6 invoked with $l = l_i, b = b_i$ and $\mathcal{T}' = \mathcal{T}'_{x,i}$, we conclude that, when V' sends i in step 4, over the randomness of $\{f_j\}, \{z_j\}$, a malicious prover fails to find $s \in \{0, 1\}^R$ such that $(Q, \text{ver}, \pi|_Q) \leftarrow \text{Sim}(x; s)$ that satisfies $\pi|_{\hat{Q}_i} \neq \pi'|_{\hat{Q}_i}$ for all $\pi' \in \text{Sup}(P(x))$ except with probability at most $1/16^3$. The lemma now follows from the claims and the observations made in the beginning of this proof.

3.4 Dealing with Private Randomness of Verifier

We address the loss in generality due to having considered only PCP schemes in which the verifier circuit is deterministic. In the more general definition, the prover samples the proof $\pi \stackrel{\$}{\leftarrow} P(x)$ on input $x \in L_n$, and verifier samples query set and verifier circuit $(Q, \text{ver}) \stackrel{\$}{\leftarrow} V(n)$ where ver is a randomized function (circuit). On the instance (x, π) , the verifier outputs $\text{ver}(x, \pi|_Q, s)$, where s is private randomness freshly sampled from a randomness domain. When defining query reusability, the verifier is allowed to sample fresh private while computing its output on each instance even though Q and ver are reused. In the full version, we establish that the techniques and negative results we developed in the previous sections can be generalized to this setting with a few careful modifications. The core idea in such a transformation is to repeat the same instance sufficiently many times so the adversary can estimate probability with which the query reusing verifier accepts the instance.

3.5 A Stronger Attack Against Reusable PCP

We construct a more involved attack A that uses only $\tilde{O}(\sqrt{q(n)})$ epochs, and provides the same set of guarantees when it fails against a query reusing PCP scheme. The new attack learns $\tilde{O}(r(n))$ new coordinates in the query set or finds and discards a subset of coordinates of size roughly $1/r(n)$ of the proof as uninfluential in every other epoch, unless the soundness is altogether broken in that epoch.

Suppose the verifier reuses $(Q, \text{ver}) \stackrel{\$}{\leftarrow} V(1^n)$ across all epochs. In any epoch, let View be A 's view, $\hat{Q}$ be the subset of Q that A has learned, and DC (as in 'don't care') be the subset A has deemed uninfluential. A partitions $[\ell] \setminus (\hat{Q} \cup \text{DC})$ into $\tilde{r} = r(n)/\log^2 n$ equal sized subsets $U_1, \dots, U_{\tilde{r}}$. For each U_i , A non-adaptively chooses a set of instances such that A can learn a coordinate in $Q \cap U_i$ from the verifier's outputs unless U_i has little influence on verifier's output in the residual PCP. For this, A chooses $x \in L_n$, π' and π'' such that both π' and π'' are 0 on all coordinates in DC and $\text{DC} \cup U_i$, respectively; agrees with a valid proof π on the remaining coordinates; and $\text{ver}'(x, \pi'|_{Q'}) = 1$ and $\text{ver}'(x, \pi''|_{Q'}) = 0$ with substantial probability when $(Q', \text{ver}') \stackrel{\$}{\leftarrow} V_{\text{view}}$. If verifier accepts (x, π') and rejects (x, π'') , A can learn a new coordinate in $U_i \cap Q$ using the same trick we used in the previous attack. Whereas, if sufficiently many such tests fail (which can be issued in parallel), A can be sure that $\text{DC} \cup U_i$ does not have any influence on the decision of the verifier. In this case, DC can be updated by adding U_i to it, and A can proceed to the next epoch.

In the above sketch, we have hidden a few of the details. Firstly, A always ensures that the residual PCP has good soundness by ensuring that any $x \in \{0, 1\}^n \setminus L_n$ and purported proof π will be rejected by the residual PCP. Further, to ensure the coordinates in DC has little influence, A carries out a much stronger test in the real attack. We provide a formal construction and proof in the full version.

The efficient attack provides the same set of guarantees provided by the attack presented in Sect. 3.1; i.e., it satisfies the statements in Lemma 3 but for $r(n) = \tilde{O}(\sqrt{q(n)})$. Hence, all the previous negative results can be obtained using the new attack with less number of epochs.

Theorem 7 (Improved Query Complexity Lower Bound). *Let (P, V) be a $(\ell(n), q(n))$ -PCP system for a language $L = \{L_n\}$. There exists an adversary that wins the $\tilde{O}(\sqrt{q(n)})$ -epoch soundness game, if L does not belong to NP/poly with $q(n)$ -bit long witness. Whereas, if the PCP scheme is statistical zero-knowledge, such an adversary exists when L is NP-complete if coNP is not contained in NP/poly.*

4 Positive Results

We move on to the positive results in the paper, which are of two kinds. The first kind are PCP schemes and statistical ZK-PCP schemes with bounded query reusability. We then construct fully query reusable statistical zero-knowledge probabilistically checkable arguments, or simply statistical ZK-PCP arguments and computational zero-knowledge proofs, or simply computational ZK-PCP under standard cryptographic assumptions.

We build a query reusing test for zero vector testing which will be a crucial tool in all the constructions with bounded query reusability.

4.1 A Zero Test with Bounded Reusable Randomness

Consider the property testing problem in which an n -bit string x is accepted if it is the all zero string 0^n , and rejected if it has Hamming weight more than $n/2$, denoted as $\|x\| \geq n/2$. We now formally describe the zero-test for n -bit strings that is query reusable for $r(n)$ epochs. The construction follows the intuition outlined in the technical overview.

For any $r(n) = n^{1/2-\delta}$ ($\delta > 0$), consider the algorithm V defined below: Let $\rho = r^2(n) \log^2 n$. $(Q = (Q_1, \dots, Q_{r(n)+1}), \text{ver}) \xleftarrow{\$} V$, where $Q_i \xleftarrow{\$} \binom{[n]}{\rho}$ for all $1 \leq i \leq r(n) + 1$, and, for any $x \in \{0, 1\}^n$,

$$\text{ver}(x|_{Q_1 \cup \dots \cup Q_{r(n)+1}}) = \begin{cases} 1 & \text{if } \|x|_{Q_i}\| \leq \frac{(i-1)\rho}{4r(n)}, \forall i \in [r(n) + 1], \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

We will show that V can reuse its randomness for $r(n)$ -epochs. Formally, for any polynomial $p(n)$, there exists a negligible function $\text{negl}(n)$ such that, for any $n \in \mathbb{N}$, any computationally unbounded adversary A that interacts with V for $r(n)$ epochs by making $p(n)$ queries in each epoch wins the following game with probability at most $\text{negl}(n)$:

1. $(Q, \text{ver}) \xleftarrow{\$} V(n)$.
2. For $\text{ep} = 1, \dots, r(n)$:

- (a) A chooses $x^{(1)}, \dots, x^{(p(n))}$ where each $x^{(i)} \in \{0, 1\}^n$.
 - (b) A receives $b^{(i)} = \text{ver}(x^{(i)}|_Q)$ for each $i \in [p(n)]$.
3. A outputs $x^* \in \{0, 1\}^n$.

A wins the game if $\|x\| \geq n/2$, and $\text{ver}(x^*|_Q) = 1$.

Theorem 8. For any $r(n) = n^{1/3-\delta}$ ($\delta > 0$), polynomial $p(n)$, there exists a negligible function $\text{negl}(n)$ such that, for any $n \in \mathbb{N}$, any computationally unbounded adversary A wins the above game with probability at most $\text{negl}(n)$.

Proof. Fix polynomials $p(n)$, $r(n)$, and an adversary A . The proof uses a hybrid argument. Fix n . We will suppress n in the following notations for simplicity. For $k = 0, \dots, r$, define Game k as follows:

1. For $1 \leq i \leq r+1$, sample $Q_i \xleftarrow{\$} \binom{[n]}{\rho}$ where $\rho = r^2(n) \log^2 n$.
2. For $1 \leq \text{ep} \leq r$:
 - (a) A chooses $\{x^{(i)}\}_{i \in [p(n)]}$, where $x^{(i)} \in \{0, 1\}^n$ for all i .
 - (b) If $\text{ep} \leq k$, A receives $\{b^{(i)} = \text{ver}'(x^{(i)}|_{Q_1 \cup \dots \cup Q_k})\}_{i \in [p(n)]}$; otherwise, A receives $\{b^{(i)} = \text{ver}(x^{(i)}|_{Q_1 \cup \dots \cup Q_{r+1}})\}_{i \in [p(n)]}$. Here, for any x ,

$$\text{ver}'(x|_{Q_1 \cup \dots \cup Q_k}) = \begin{cases} 1 & \text{if } \|x|_{Q_i}\| \leq \frac{(i-1)\rho}{4r(n)}, \forall i \in [k], \\ 0 & \text{otherwise.} \end{cases}$$

3. A outputs x^* .

A wins if $\|x^*\| \geq n/2$ and $\text{ver}(x^*|_{Q_1 \cup \dots \cup Q_{r+1}}) = 1$.

Observe that Game 0 is identical to the r -sound soundness game used in the theorem statement. We prove the following claim.

Claim. For any $k = 1, \dots, r+1$, the winning probability of A in Game $k-1$ and Game k are apart by a value that is negligible in n .

Proof. To prove the claim, it suffices to show that the joint distribution of $(x^{(i)}, b^{(i)})_{i \in [p]}$ in each epoch of game $k-1$ and game k are at most negligibly far in total variation distance. A receives $b^{(i)} = \text{ver}(x^{(i)}|_{Q_1}, \dots, x^{(i)}|_{Q_{r+1}})$ for all $i \in [p]$ in epoch k of Game $k-1$. Whereas, in epoch k of Game k , A receives $b^{(i)} = \text{ver}'(x^{(i)}|_{Q_1}, \dots, x^{(i)}|_k)$ for all $i \in [p]$. This is the only difference between the two games. To prove the claim, it suffices to prove that, for any fixing of $Q_1, \dots, Q_{k-1}$ as subsets of $[n]$ of size $\rho = r^2 \log^2(n)$, and for any computationally unbounded adversary B^* ,

$$\Pr \left[\begin{array}{c} \exists j \in [p], \text{ver}(\{x^{(k,j)}\}_{\ell \in [r+1]}) \\ \neq \text{ver}'(\{x^{(k,j)}\}_{\ell \in [k]}) \end{array} \middle| \begin{array}{c} Q_j \leftarrow \binom{[n]}{\rho}, k \leq j \leq r+1, \\ (x^{(k,1)}, \dots, x^{(k,p)}) \leftarrow B^* \end{array} \right] \leq \text{negl}(n). \quad (12)$$

For any $x \in \{0, 1\}^n$ and any fixed $Q_1, \dots, Q_{k-1}$, the event E described as $\text{ver}(\{x|_{Q_\ell}\}_{\ell \in [r+1]}) \neq \text{ver}'(\{x|_{Q_\ell}\}_{\ell \in [k]})$ occurs only if $\|x|_{Q_k}\| \leq \epsilon_k$ and $\|x|_{Q'_k}\| >$

$\epsilon_{k'}$ for some $k < k' \leq p$ when $\epsilon_{k''}$ is defined to be $\frac{(k''-1)\rho}{4r(n)}$, for all $k'' \in [p]$. We will argue that E occurs with negligible probability when $Q_k, \dots, Q_{r+1}$ are drawn from $[n] \setminus (Q_1 \cup \dots \cup A_{k-1})$, and noting that $|Q_1 \cup \dots \cup Q_{k-1}| \leq n/2$. The actual case can be analyzed along similar lines with a few small changes. We have

$$\begin{aligned} \Pr[E] &\leq \sum_{k'=k+1}^{r+1} \Pr[(\|x|_{Q_k}\| \leq \epsilon_k) \wedge (\|x|_{Q_{k'}}\| > \epsilon_{k'})] \\ &\leq \sum_{k'=k+1}^{r+1} 2e^{-\rho(\epsilon_{k'} - \epsilon_k)^2 / 2(\epsilon_k + \epsilon_{k'})} \\ &\leq r2e^{-r^2 \log^2 n \cdot (\frac{\epsilon}{2r})^2 / 8r} = O(2^{-\log^2 n}). \end{aligned}$$

Above, the first inequality used a union bound, second inequality used Lemma 1, and third inequality used the facts $\epsilon_k - \epsilon_{k'} \geq \epsilon/2r$ for any $k' > k$ and $\epsilon_k + \epsilon_{k'} \leq 2(r+1) \leq 4r$ when $r \geq 1$. Equation (12) now follows by a union bound. The proof is complete.

Claim. In Game r , the winning probability of A is negligible in n .

Proof. In Game r , Q_{r+1} may be equivalently sampled after A chooses x^* . But then, for any x^* of hamming weight more than $n/2$, over the randomness of Q_{r+1} , by a Chernoff bound,

$$\Pr[\|x^*|_{Q_{r+1}}\| \leq \rho/4] \leq \text{negl}(n).$$

The claim follows.

The theorem follows from the above two claims by a triangular inequality.

4.2 PCP with Bounded Reusability

Theorem 9. Suppose (P, V) is a $(\ell(n), q(n))$ -PCP. For any polynomial $r(n)$, there exists a $(\ell(n), q'(n))$ -PCP (P, V') that is query reusable for $r(n)$ -epochs when $q'(n) = r^3(n) \cdot q(n) \cdot \log^2 n$.

Proof. Fix $n \in \mathbb{N}$. Let $\{0, 1\}^R$ be the randomness used by $V(1^n)$. That is, the distribution $V(1^n) = V(1^n; r)$ where $s \xleftarrow{\$} \{0, 1\}^R$. Define $V'(1^n)$ as follows (suppressing n in notations):

PCP with $r(n)$ -epoch query reusability

1. For $i \in [r]$ and $j \in [r^2 \log^2 n]$: sample $s^{(i,j)} \xleftarrow{\$} \{0, 1\}^R$. For each i, j , let $(Q^{(i,j)}, \text{ver}^{(i,j)}) \leftarrow V(1^n; s^{(i,j)})$.
2. Define $V'(1^n, \{s^{(i,j)}\}) = (Q', \text{ver}')$ where $Q' = \{Q^{(i,j)}\}_{i,j}$ and circuit ver' take inputs $x \in \{0, 1\}^n$ and $\{\pi|_{Q^{(i,j)}}\}_{i,j}$ and compute $b^{(i,j)} \leftarrow \text{ver}^{(i,j)}(x, \pi|_{Q^{(i,j)}})$, and outputs 1 if, number of j such that $b^{(i,j)} = 0$ is at most $\lfloor r(i-1)/4 \rfloor$ for each $i \in [r]$, else outputs 0.
3. Output (Q', ver') .

The completeness of $V'(1^n)$ follows since, when $x \in L_n$ and π is a valid proof, $b^{(i,j)} = 1$ for all i, j , which in turn follows from the completeness of $V(1^n)$. The r -epoch soundness is argued by a reduction from r -epoch soundness of the zero test which accepts $z \in \{0, 1\}^{2^R}$ if z is the all zero string, and rejects if the Hamming weight of z is more than 2^{R-1} .

Given a computationally unbounded adversary A that wins the $r(n)$ -epoch soundness game against $V'(1^n)$ with probability μ , we will show that adversary B wins the following game with probability μ as follows:

1. For $i \in [r]$ and $j \in [r^2 \log^2 n]$: sample $s^{(i,j)} \leftarrow \{0, 1\}^R$.
2. Let $(\tilde{Q}, \widetilde{\text{ver}}) \leftarrow \tilde{V}(1^n; \{s^{(i,j)}\})$ where $\tilde{Q} = \{Q_i\}$ and $Q_i = \{s^{(i,j)}\}_j$ for each $i \in r$. $\widetilde{\text{ver}}$ is such that $\widetilde{\text{ver}}(z|_{\tilde{Q}})$ is 1 if the Hamming weight of $z|_{Q_i}$ is at most $\lfloor r(i-1)/4 \rfloor$ for each $i \in [r]$, and 0 otherwise.
3. For $i = 1, \dots, r$:
 - (a) B receives $\{(x^{(i,j)}, \pi^{(i,j)})\}_{j \in [p]}$ from A .
 - (b) For each $j \in [p]$, B computes $z^{(i,j)} \leftarrow \phi(x^{(i,j)}, \pi^{(i,j)})$, where ϕ maps any (x, π) to z such that, for each $s \in \{0, 1\}^{2^R}$

$$z_s = 1 - \text{ver}(x, \pi|_Q), \text{ where } (Q, \text{ver}) \leftarrow V(1^n; s).$$
 - (c) B uses $\{z^{(i,j)}\}_{j \in [p]}$ as inputs to receive $b^{(i,j)} \leftarrow \widetilde{\text{ver}}(z|_{\tilde{Q}})$ for each j . B forwards $\{b^{(i,j)}\}_{i,j}$ to A .
4. At any point, if A outputs (x, π) , B outputs $\phi(x, \pi)$.

We will now argue that B wins the above game with probability μ .

Claim. For any $\{s^{(i,j)}\}_{i \in [r], j \in [r^2 \log^2 n]}$ where each $s^{(i,j)} \in \{0, 1\}^R$, and (x, π) , when $(Q', \text{ver}') \leftarrow V'(1^n; \{s^{(i,j)}\})$ and $(\tilde{Q}, \widetilde{\text{ver}}) \leftarrow \tilde{V}(1^n; \{s^{(i,j)}\})$,

$$\text{ver}'(x, \pi|_{Q'}) = \widetilde{\text{ver}}(\phi(x, \pi)|_{\tilde{Q}}).$$

The claim follows from the definition of ϕ .

Observe that both $\tilde{V}$ and V' sample $\{s^{(i,j)}\}_{i \in [r], j \in [r^2 \log^2 n]}$ according to the same distribution. Hence, in the above game, the responses $b^{(i,j)}$ forwarded by B coincides with the responses A would get when it interacts with $V'(1^n; \{s^{(i,j)}\})$, which occurs with the same probability. If A outputs (x, π) which wins the r -epoch soundness game against V' , then $x \in \{0, 1\}^n \setminus L_n$ and $\text{ver}'(x, \pi|_{Q'}, \text{st}') = 1$. But then, by the above claim, $z = \phi(x, \pi)$ is such that $\widetilde{\text{ver}}(z) = 1$. But, since $x \in \{0, 1\}^n \setminus L_n$, $\text{ver}(x, \pi|_Q) = 1$ with probability less than $1/2$ (in fact this probability is negligible) when $(Q, \text{ver}) \leftarrow V(1^n)$. But then, the Hamming weight of z is at least 2^{R-1} (recall $z_r = 1$ if $\text{ver}(x, \pi|_Q) = 1$ when $(Q, \text{ver}) \leftarrow V(n; r)$.) This contradicts the r -epoch soundness of the zero test as established in Theorem 8 whenever $2^R \geq r^3$, since B has essentially broken the r -epoch soundness of the zero test.

4.3 Statistical ZK-PCP Proofs with Bounded Reusability

We define the notion of $r(n)$ -epoch zero-knowledge below.

Definition 5 (Query Reusable ZK-PCP). Let $L = \{L_n\}_{n \in \mathbb{N}}$ be a language in NP. (P, V) constitute a query reusable $(\ell(n), q(n))$ -ZK-PCP, if it is a query reusable $(\ell(n), q(n))$ -PCP for L which additionally satisfies $r(n)$ -epoch zero-knowledge property:

$r(n)$ -Epoch Zero Knowledge. There exists a polynomial time randomized simulator Sim such that for any n and $x_1, \dots, x_{r(n)} \in L_n$, consider the following two ensembles.

$$\left(Q, \text{ver}, \pi_1|_Q, \dots, \pi_{r(n)}|_Q \mid (Q, \text{ver}, \pi_1|_Q, \dots, \pi_{r(n)}|_Q) \xleftarrow{\$} \text{Sim}(x_1, \dots, x_{r(n)}) \right), \\ \left(Q, \text{ver}, \pi_1|_Q, \dots, \pi_{r(n)}|_Q \mid \pi_j \xleftarrow{\$} P(x_j), \forall j \in [r(n)], (Q, \text{ver}) \xleftarrow{\$} V(1^n) \right).$$

If the two ensembles are perfectly/statistically/computationally indistinguishable then (P, V) is perfect/statistical/computational ZK-PCP.¹

Theorem 10. Suppose (P, V) is a $(\ell(n), q(n))$ -ZK-PCP. For any polynomial $r(n)$, there exists a $(\ell'(n), q'(n))$ -ZK-PCP (P', V') scheme that is reusable for $r(n)$ -epochs when $\ell'(n) = r^{(n)} \cdot \ell(n)$, $q'(n) = r^3(n) \cdot q(n) \cdot \log^2 n$.

Proof. The construction of (P', V') from (P, V) is similar to that of the bounded randomness reusable PCP scheme in Theorem 9, with one crucial difference. In Theorem 9, the new PCP scheme was built by simply sampling several independent copies of the given verifier and applying statistical tests to r subsets of these copies with progressively larger thresholds. In a reusable ZK-PCP, the verifier can not query the same proof multiple times without breaking zero-knowledge. To overcome this, P' publishes $r^3(n) \cdot \log^2 n$ independent (randomly sampled) proofs so that the verifier may query each copy only once. This preserves the zero-knowledge property of a single execution of the reusable ZK-PCP since $r^3(n) \cdot \log^2 n$ is still polynomial in n . This construction, however, cannot be proved to be zero-knowledge when proving multiple statements while reusing the query. This is because, in general, the simulator for the underlying ZK-PCP scheme outputs a fresh query set and verifier circuit in every invocation.

We get around this by using ZK-PCP for semi-malicious verifiers. In detail, the MPC-based constructions of ZK-PCP [15, 17, 21] carry out simulation by invoking the MPC-based simulator, where the PCP queries correspond to the set of corrupted parties. Consequently, these constructions provide ZK against semi-malicious verifier since the simulator can be invoked with (Q, ver) in addition to x . By instantiating our construction with such a ZK-PCP, we can simulate the query reusing verifier's view while proving polynomially many statements.

The analysis of $r(n)$ -epoch soundness is similar to the one in the proof of Theorem 9.

¹ A stronger notion of zero-knowledge requires the private randomness used by V to generate the reused (Q, ver) to be jointly simulated with all other random variables. Our construction indeed satisfies this requirement.

4.4 Fully Query Reusable Computational ZK-PCP

We realize fully query reusable computational ZK-PCP proof and statistical ZK-PCP argument.

The starting point for our construction is non-interactive zero-knowledge proofs/arguments (NIZK) in the common reference string (CRS) model. Since NIZK allow reuse of the CRS, they directly imply fully reusable ZK-PCP in the CRS model. When the verifier is semi-honest, the CRS can be sampled once and for all by the verifier, and the prover can use it to generate all the proofs. To realize this non-interactively, consider the functionality taking input x , witness w and randomness s from the prover and the CRS u from the verifier, and delivers the proof computed as a function of x, w, s, u to the verifier. We will argue that implementing this functionality using secure Non-Interactive Secure Computation protocol with statistical semi-honest security in the OT-hybrid model (NISC/OT) [16] implies a honest verifier ZK-PCP.

A NISC/OT protocol securely realizes the above functionality using only parallel calls to ideal OT functionality. Let π be the concatenation of prover's inputs to the parallel OT calls (based on their inputs (x, w, s) to the functionality) in the NISC/OT protocol realizing the above functionality. The prover publishes π as the ZK-PCP for instance x . An honest verifier computes its input to the parallel OT calls, based on their input u to the functionality, and picks up the appropriate coordinates of π as dictated by their OT input. The 'OT outputs' can be used to recover the NIZK proof computed by the underlying functionality; verifier can now use this proof to accept/reject the input by executing the verifier algorithm of the NIZK. It is easy to see that the scheme inherits its ZK property from the underlying NIZK since prover uses fresh randomness for every proof. More crucially, a malicious prover can simulate the output of the NISC/OT computation using verifier's input u which is essentially the CRS. If the prover breaks the soundness of the ZK-PCP scheme, it directly implies a proof that breaks the underlying NIZK scheme. This shows that the ZK-PCP is a fully reusable argument with statistical ZK (respectively, fully reusable proof with computational ZK) if the underlying NIZK is an argument with statistical ZK (respectively, proof with computational ZK).

Efficient, statistical UC secure NISC/OT protocols exist for any function in the NC^1 complexity class [16]. In [25], Waters constructed a computational NIZK proof and a statistical NIZK argument under learning with errors assumption. The functionality implied by this NIZK cannot be computed using an NC^1 circuit; however a modified functionality that takes a preprocessing of prover's inputs x, w, s and the input u of the verifier can be computed using an NC^1 circuit. Using this functionality we obtain the desired result.

Theorem 11. *Assuming learning with errors (with subexponential modulus to noise ratio), there exists fully reusable statistical ZK-PCP arguments and computational ZK-PCP proofs.*

Acknowledgements. We thank Omri Ben-Eliezer, Moni Naor, and Eylon Yogev for helpful discussions and pointers, and the anonymous Eurocrypt reviewers for their

helpful comments. Y. Ishai and E. Kushilevitz were supported by ISF grants 2774/20 and 3527/24 and BSF grant 2022370. Y. Ishai was additionally supported by ISF-NSFC grant 3127/23. R. Ostrovsky and V. Narayanan were supported in part by NSF grants CNS-2246355, CCF-2220450, US-Israel BSF grant 2022370, and by Sunday Group.

References

1. Ames, S., Hazay, C., Ishai, Y., Venkatasubramanian, M.: Ligerio: lightweight sub-linear arguments without a trusted setup. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 2087–2104. ACM Press, October/November 2017
2. Arora, S., Safra, S.: Probabilistic checking of proofs; a new characterization of NP. In: 33rd FOCS, pp. 2–13. IEEE Computer Society Press, October 1992
3. Barak, B., Goldreich, O.: Universal arguments and their applications. In: Proceedings 17th IEEE Annual Conference on Computational Complexity, pp. 194–203. IEEE (2002)
4. Bitansky, N., et al.: The hunting of the SNARK. *J. Cryptol.* **30**(4), 989–1066 (2017)
5. Lawrence Carter, J., Wegman, M.N.: Universal classes of hash functions. *J. Comput. Syst. Sci.* **18**(2), 143–154 (1979)
6. Chase, M., et al.: Post-quantum zero-knowledge and signatures from symmetric-key primitives. In: Thuraisingham, B.M., Evans, D., Malkin, T., Xu, D. (eds.) ACM CCS 2017, pp. 1825–1842. ACM Press, October/November 2017
7. Chase, M., et al.: Reusable non-interactive secure computation. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11694, pp. 462–488. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_15
8. Choudhuri, A.R., Jain, A., Jin, Z.: Non-interactive batch arguments for NP from standard assumptions. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021. LNCS, vol. 12828, pp. 394–423. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-84259-8_14
9. Couteau, G., Jain, A., Jin, Z., Quach, W.: A note on non-interactive zero-knowledge from CDH. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part IV. LNCS, vol. 14084, pp. 731–764. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-38551-3_23
10. Feige, U., Lapidot, D., Shamir, A.: Multiple non-interactive zero knowledge proofs based on a single random string (extended abstract). In: 31st FOCS, pp. 308–317. IEEE Computer Society Press, October 1990
11. Fortnow, L.: The complexity of perfect zero-knowledge (extended abstract). In: Aho, A. (ed.) 19th ACM STOC, pp. 204–209. ACM Press, May 1987
12. Giacomelli, I., Madsen, J., Orlandi, C.: Zkboo: faster zero-knowledge for Boolean circuits. In: Holz, T., Savage, S. (eds.) 25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, 10–12 August 2016, pp. 1069–1083. USENIX Association (2016)
13. Goldwasser, S., Sipser, M.: Private coins versus public coins in interactive proof systems. In: 18th ACM STOC, pp. 59–68. ACM Press, May 1986
14. Gur, T., O’Connor, J., Spooner, N.: A zero-knowledge PCP theorem. *CoRR*, abs/2411.07972 (2024)
15. Ishai, Y., Kushilevitz, E., Ostrovsky, R., Sahai, A.: Zero-knowledge from secure multiparty computation. In: Johnson, D.S., Feige, U. (eds.) 39th ACM STOC, pp. 21–30. ACM Press, June 2007

16. Ishai, Y., Prabhakaran, M., Sahai, A.: Founding cryptography on oblivious transfer – efficiently. In: Wagner, D. (ed.) CRYPTO 2008. LNCS, vol. 5157, pp. 572–591. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_32
17. Ishai, Y., Weiss, M.: Probabilistically checkable proofs of proximity with zero-knowledge. In: Lindell, Y. (ed.) TCC 2014. LNCS, vol. 8349, pp. 121–145. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54242-8_6
18. Jain, A., Jin, Z.: Non-interactive zero knowledge from sub-exponential DDH. In: Canteaut, A., Standaert, F.-X. (eds.) EUROCRYPT 2021. LNCS, vol. 12696, pp. 3–32. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77870-5_1
19. Kilian, J.: A note on efficient zero-knowledge proofs and arguments. In: Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing, pp. 723–732 (1992)
20. Kilian, J., Micali, S., Ostrovsky, R.: Minimum resource zero-knowledge proofs (extended abstract). In: 30th FOCS, pp. 474–479. IEEE Computer Society Press, October/November 1989
21. Kilian, J., Petrank, E., Tardos, G.: Probabilistically checkable proofs with zero knowledge. In: 29th ACM STOC, pp. 496–505. ACM Press, May 1997
22. Lombardi, A., Quach, W., Rothblum, R.D., Wichs, D., Wu, D.J.: New constructions of reusable designated-verifier NIZKs. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11694, pp. 670–700. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_22
23. Micali, S.: A secure and efficient digital signature algorithm. Technical Memo MIT/LCS/TM-501b, Massachusetts Institute of Technology, Laboratory for Computer Science, April 1994
24. Peikert, C., Shiehian, S.: Noninteractive zero knowledge for NP from (plain) learning with errors. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11692, pp. 89–114. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26948-7_4
25. Waters, B.: A new approach for non-interactive zero-knowledge from learning with errors. In: Mohar, B., Shinkar, I., O’Donnell, R. (eds.) Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, 24–28 June 2024, pp. 399–410. ACM (2024)